

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

This Page Blank (uspto)



Europäisches
Patentamt

European
Patent Office

Office eur péen
des brevets

#3 9-8-01
Priority Papers



Bescheinigung

Certificate

Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

00116861.6

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

I.L.C. HATTEN-HECKMAN

DEN HAAG, DEN
THE HAGUE, 15/03/01
LA HAYE, LE

This Page Blank (uspto)



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

Blatt 2 der Bescheinigung
Sheet 2 of the certificate
Page 2 de l'attestation

Anmeldung Nr.:
Application no.:
Demande n°: 00116861.6

Anmeldetag:
Date of filing: 04/08/00
Date de dépôt:

Anmelder:
Applicant(s):
Demandeur(s):
International Business Machines Corporation
Armonk, NY 10504
UNITED STATES OF AMERICA

Bezeichnung der Erfindung:
Title of the invention:
Titre de l'invention:

Method and system for the automatic segmentation of an audio stream into semantic or syntactic units

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s) revendiquée(s)

Staat:
State:
Pays:

Tag:
Date:
Date:

Aktenzeichen:
File no.
Numéro de dépôt:

Internationale Patentklassifikation:
International Patent classification:
Classification internationale des brevets:

/

Am Anmeldetag benannte Vertragsstaaten:
Contracting states designated at date of filing: AT/BE/CH/CY/DE/DK/ES/FI/FR/GB/GR/IE/IT/LI/LU/MC/NL/PT/SE/TR
Etats contractants désignés lors du dépôt:

Bemerkungen:
Remarks:
Remarques:

This Page Blank (uspto)

D E S C R I P T I O N

METHOD AND SYSTEM FOR THE AUTOMATIC SEGMENTATION OF AN AUDIO
STREAM INTO SEMANTIC OR SYNTACTIC UNITS

BACKGROUND OF THE INVENTION

The present invention relates to computer-based automatic segmentation of audio streams like speech or music into semantic or syntactic units like sentence, section and topic units, and more specifically to a method and a system for such a segmentation wherein the audio stream is provided in a digitized format.

Prior to or during production of audiovisual media like movies or broadcasted news a huge amount of raw audio material is recorded. This material is almost never used as recorded but subjected to editing, i.e. segments of the raw material relevant to the production are selected and assembled into a new sequence.

Today this editing of the raw material is a laborous and time-consuming process involving many different steps, most of which require human intervention. A crucial manual step during the preprocessing of the raw material is the selection of cut points, i.e. the selection of possible segment boundaries, that may be used to switch between different segments of the raw material during the subsequent editing steps. Currently these cut points are selected interactively in a process that requires a human to listen for potential audio cut points like speaker changes or the end of a sentence, to determine the exact time at which the cut point occurs, and to add this time to an edit- decision list (EDL).

For the above reasons there is a substantial need to automate part or even all of the above preprocessing steps.

Thereupon many audio or audiovisual media sources like real media streams or recordings of interviews, conversations or news broadcasts are available only in audio or audiovisual form i.e. they lack in corresponding textual transcripts and thus in typographic cues such as headers, paragraphs, sentence punctuation, and capitalization which would allow for segmentation of those media streams only by linking transcript information to audio (or video) information as proposed in European Patent Application X XXX XXX (docketno. DE9-1999-0053 of present applicant). But those cues are absent or hidden in speech output.

Therefore a crucial step for the automatic segmentation of such media streams is (automatic) determination of semantic or syntactic boundaries like topics, sentences and phrase boundaries.

There exist approaches which use prosody whereby prosodic features are those features which have not only influence on single media stream segments (called phonemes) but extend over a number of segments. Exemplary prosodic features are information extracted from the timing and melody of speech like pausing, changes in pitch range or amplitude, global pitch declination, or melody and boundary tone distribution for the segmentation process.

As disclosed in a recent article by E. Shriberg, A. Stolcke et al. entitled „Prosody-Based Automatic Segmentation of Speech into Sentences and Topics“, published as pre-print to appear in Speech Communication 32(1-2) in September 2000, evaluation of prosodic features usually is combined with statistical language models mainly based on Hidden Markov Theory and thus presume words already decoded by a speech recognizer. The advantage to use prosodic indicators for segmentation is that prosodic features are relatively unaffected by word identity and thus improve the robustness of the entire segmentation process.

A further article by A. Stolcke, E. Shriberg et al. entitled „Automatic Detection of Sentence Boundaries and Disfluencies Based on Recognized Words“, published as Proceedings of the International Conference on Spoken Language Processing, Sydney, 1998, concerns also segmentation of audio streams using a prosodic model and is accordingly based on speech already transcribed by an automatic speech recognizer.

In the above at first cited article, prosodic modeling is mainly based on only very local features whereby for each inter-word boundary prosodic features of the word immediately preceding and following the boundary, or alternatively within an empirically optimized window of 20 frames before and after the boundary, are analyzed. In particular, prosodic features are extracted which reflect pause durations, phone durations, pitch information, and voice quality information. Pause features are extracted at the inter-word boundaries. Pause duration, a fundamental frequency (F0), and voice quality features are extracted mainly from the word and window preceding the boundary. In addition, pitch-related features reflecting the difference in pitch across the boundary are included in the analysis.

In the above article by E. Shriberg et al., chapter 2.1.2.3, it is generally referred to a mechanism for determining the F0 signal. A similar mechanism according to the invention which will be discussed in more detail later referring to Fig. 3. The further details of the F0 processing are of no relevance for the understanding of the present invention.

As mentioned above further, the segmentation process disclosed in the precited article is based on language modeling in order to capture information about segment boundaries contained in the word sequences. Therein described approach is to model the joint distribution of boundary types and words in a Hidden Markov Model (HMM), the hidden variable being the word boundaries. For the segmentation of sentences it is therefore relied on a hidden-event N-gram language model where the states of the HMM

consist of the end-of-sentence status of each word, i.e. boundary or no-boundary, plus any preceding words and possible boundary tags to fill up the N-gram context. As commonly known in the related art, transition probabilities are given by N-gram probabilities which in this case are estimated from annotated, boundary-tagged user-specific training data.

Concerning segmentation, the authors of that article further propose use of a decision tree where a number of prosodic features are used which fall into different groups. In order to designate the relative importance of these features in the decision tree, a measure called „feature usage“ is utilized which is computed as the relative frequency with which that feature or feature class is queried in the decision tree. As prosodic features are used pause duration at a given boundary, turn/no turn at the boundary, F0 difference across the boundary, and rhyme duration.

The above cited prior art approaches have the drawback that they either necessarily use a speech recognizer or that they require multiple processing steps. Therupon the entire segmentation process is error-prone, i.e. has to rely on speech recognizer output, and is time-consuming. In addition, most of the known approaches use rather complex mechanisms and technologies and thus their technical realization is rather cost-extensive.

SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide a method and an apparatus for segmentation of audio streams which perform the segmentation as automatic as possible.

A further object is to provide such a method and apparatus which can be implemented with minimum technical efforts and requirements thus being minimum cost-extensive.

It is another object to provide such a method and apparatus which perform the segmentation of audio streams as robust as possible.

It is another object to provide such a method and system which allow for segmentation of a continuous audio stream without the provision of any corresponding transcript.

Another object is to provide an automatic, real-time system for the predescribed segmentation of audio streams.

Yet another object is to provide such a method and apparatus which allow for a not user-specific segmentation of audio streams.

The objects are solved by the features of the independent claims. Advantageous embodiments of the invention are subject matter of the dependent claims.

The invention accomplishes the foregoing by determining a fundamental frequency for the digitized audio stream, detecting changes of the fundamental frequency in the audio stream, determining candidate boundaries for the semantic or syntactic units depending on the detected changes of the fundamental frequency, extracting at least one prosodic feature in the neighborhood of the candidate boundaries, and determining boundaries for the semantic or syntactic units depending on the at least one prosodic feature.

The idea or concept underlying the present invention is to provide a pre-segmentation of the audio stream and thereby obtaining potential or candidate boundaries between semantic or syntactic units, preferably based on the fundamental frequency F_0 . The concept is based on the observation that sonorant i.e. voiced audio segments, which are characterized by $F_0 = 0N$ according to the invention, only rarely extend over two semantic or syntactic units like sentences.

It is emphasized hereby that the candidate boundaries are obtained by applying only one criterion namely whether F0 is ON or OFF.

Based on the obtained candidate boundaries, prosodic features are extracted at these boundaries, preferably in both (time) directions starting from a particular candidate boundary. In particular, continuous features relevant for prosody are used.

In a preferred embodiment of the invention, an index function is defined for the fundamental frequency having a value = 0 if the fundamental frequency is undefined and having a value = 1 if the fundamental frequency is defined. It is, among others, the so-called Harmonics-to-Noise-Ratio which allows for the predication whether the F0 is defined or undefined using a threshold value (see Boersma(1993) for details). The index function allows for an automatization of the pre-segmentation process for finding candidate boundaries and thus the following steps of processing prosodic features can be performed automatically too at the candidate boundaries.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will be understood more readily from the following detailed description when taken in conjunction with the accompanying drawings, in which:

- Fig. 1 is a piece cut out from a typical continuous audio stream which can be segmented in accordance with the invention;
- Fig. 2 depicts typical F0 data b. calculated from a real life speech signal a.;
- Fig. 3 is a block diagram of an apparatus for processing F0 according to the prior art;

Fig. 4a-c are diagrams for illustrating the method for segmentation of audio streams according to the invention;

Fig. 5a-c are flow diagrams for illustrating procedural steps of audio segmentation according to the invention; and

Fig. 6 is an exemplarily trained tree structure used for an audio segmentation process according to the invention.

DETAILED DESCRIPTION OF THE DRAWINGS

Fig. 1 shows a piece cut out from a digitized continuous speech signal. The original audio stream is digitized using known digitizing tools, e.g. wav or mpg-format generating software, and stored in a file as a continuous stream of digital data. It is emphasized hereby that the below described mechanism for segmentation of such an audio stream can be accomplished either in an off-line or a real-time environment. In addition, it can be implemented so that the different procedural steps are performed automatically.

It is also noted hereby that potential semantic or syntactic units, in which the audio stream can be segmented, are paragraphs, sentences or changes from one speaker to another speaker or even scenes in case of audio-visual streams segmented via the audio part.

In a first step (not shown here) the digitized audio stream is segmented into speech and non-speech segments by means of algorithms as known in the art, e.g. the one described in an article by Claude Montacié, entitled „A Silence/Noise/Music/Speech Splitting Algorithm“, published in Proceeding of the International Conference on Spoken Language Processing, Sydney, 1998. Hereby it is guaranteed that the

following steps will only be applied to speech segments.

In a next step, the fundamental frequency of the audio stream is continuously determined by use of a processor depicted in Fig. 3 and described later in more detail. Fig. 2a depicts a piece cut out from an audio stream as shown in Fig. 1, wherein Fig. 2b shows an F0 contour determined from the audio piece shown in Fig. 2a.

An applicable algorithm for the F0 processing is described in detail in an article by Paul Boersma (1993): "Accurate short-term analysis of the fundamental frequency and the harmonics-to-noise ratio of a sampled sound", Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam 17: 97-110.

The fundamental frequency F0 is the component of the spectrum of a periodic signal which comprises the longest period. To calculate it in the time domain an autocorrelation function is used whereby the fundamental frequency F0 can be obtained from the inverse of the signal period since the autocorrelation function is maximum for multiples of the period. In practice the audio signal is scanned by a window.

The above algorithm thus performs an acoustic periodicity detection on the basis of an accurate autocorrelation method based on cepstrum or combs, or the original autocorrelation methods. Boersma recognized the fact that if one wants to estimate a signal's short-term autocorrelation function on the basis of a windowed signal, the autocorrelation function of the windowed signal should be divided by the autocorrelation function of the window. For the further details it is referred to the above cited article which is regarded to be fully incorporated by reference.

Fig. 3 shows a block diagram of a simplified algorithm for the F0 detection in a digitized signal according to the previously cited article of P. Boersma (1993).

In a first step parameters like e.g. the frame length are initialized, a windowing function is selected and the autocorrelation function of the window is computed. Any windowing function used in signal processing applications can be used as a window, the preferred window function is a Gaussian window.

For each frame the following computations are performed: The windowed autocorrelation r_x of the frame is computed. The windowed autocorrelation is defined as the normalized autocorrelation of the windowed signal r_a divided by the normalized autocorrelation of the window r_w . Normalisation is carried out by dividing the autocorrelation by the value at lag zero. The following equations summarize the prescribed procedure:

$$r_x(\tau) = \frac{r_a(\tau)}{r_w(\tau)}$$

$$r_a(\tau) = \frac{\int_0^{T-\tau} a(t)a(t+\tau)dt}{\int_0^T a^2(t)dt}$$

where the variables are:

T duration of the frame

$a(t)$ the windowed signal shifted to mean 0

τ lag

From the windowed autocorrelation at most n (frequency, strength) coordinates of candidates for the fundamental frequency are

selected according to the following rules:

- The preferred value for n is 4
- The local maxima of the windowed autocorrelation are determined by parabolic interpolation
- The first candidate is the unvoiced
- The other candidates are the first $n-1$ maxima with the highest local strength

The strength R of the unvoiced candidate is computed according to

$$R = a + \max \left(0, 2 - \frac{(\text{local absolute peak})/b}{c/(1+a)} \right)$$

The preferred value for the constant a is a threshold for voicedness, for the constant b the maximal amplitude of the signal, and for c a threshold for silence.

The strength R of the other candidates is computed according to

$$R = r(\tau) - d^2 \log(e \cdot \tau)$$

where τ is the lag of the local maximum. The preferred value for d is 0.01, and for e the frequency minimum for the fundamental frequency.

Using well known dynamic programming algorithms, a path through the candidate fundamental frequencies is selected that minimizes the cost for voiced/unvoiced transitions and for octave jumps. The cost function is given by

$$\text{cost}(\{p_n\}) = \sum_{n=2}^{\text{numberOfFrames}} \text{transitionCost}(F_{n-1, p_{n-1}}, F_{np_n}) - \sum_{n=1}^{\text{numberOfFrames}} R_{np_n}$$

$$\text{transitionCost}(F_1, F_2) = \begin{cases} 0 & \text{if } F_1 = 0 \text{ and } F_2 = 0 \\ \text{VoicedUnvoicedCost} & \text{if } F_1 = 0 \text{ xor } F_2 = 0 \\ \text{OctaveJumpCost} \cdot \left| 2 \log \frac{F_1}{F_2} \right| & \text{if } F_1 \neq 0 \text{ and } F_2 \neq 0 \end{cases}$$

where

pn path through the candidates
 F frequency of candidate
 R strength of candidate

The preferred values for VoicedUnvoicedCost and OctaveJumpCost parameters are 0.2.

The result of the F0 processing is a pre-segmentation of the audio stream into segments where the segment boundaries comprise transitions between voiced and unvoiced audio sections.

The extraction of prosodic features is illustrated now referring to Figures 4a - 4c.

Fig. 4a shows a plot of F0 (in units of Hertz) over time (in units of milliseconds). The sampling rate is 10 milliseconds and thus the time distances between the plotted data is also 10 msec so far as they are not interrupted by voiceless sections. Near its center, between 34000 and 35000 msec, the plot comprises such a voiceless section which, in the underlying audio stream, separates two sentences. At the bottom of the diagram, an index function comprising values 1 (= ON) and 0 (= OFF) is depicted which is ON for voiced sections and OFF for voiceless sections of the audio stream.

The plot diagram shown in Fig. 4b, in addition to the F0 data, depicts intensity data (smaller dots) in units of decibel. The extraction of prosodic features is generally performed in the environment of voiceless sections where F0 is OFF, as the depicted section between 34000 and 35000 msec which have a time duration larger than a threshold value, e.g. 100 msec. A first feature is the length w of the voiceless section which strongly correlates with the length of pauses. In other words, the longer w is the more likely the section comprises a pause of the same length w . The features $f1$ and $f2$ represent the F0 values at the boundaries of the voiceless section w called Offset $f1$ and Onset $f2$ i.e. $f1$ is the F0 offset before the voiceless section w and $f2$ is the F0 onset after the voiceless section w . The Offset $f1$ is of greater importance than the Onset $f2$ since it has been found that F0 at the end of an audio segment of spoken language in most cases is lower than the average value.

A further prosodic feature is the difference $f2 - f1$. It has been found that a speaker after a segment boundary in most cases does

not continue with the same pitch thus resulting in a so-called pitch reset.

Another prosodic feature is the arithmetic mean m_i of the signal intensity within the voiceless section w_i . Using that feature it is possible to distinguish consecutive unvoiced sounds from real pauses.

Besides the above features, the following prosodic features can be extracted from the slope of the F_0 contour depicted in Fig. 4a and 4b which is illustrated by reference to Fig. 4c. Hereby the declination within semantic units like sentences and phrases is used to detect semantic unit boundaries or ends. The problem is that it is rather difficult to extract this feature information from F_0 out of a continuous audio stream since a standard linear regression can be performed only if the boundaries are already known. The present invention solves this by performing a linear regression only in the voiced sections directly preceding or succeeding the voiceless section w_i and only within a predetermined time window, i.e. in the present case starting from 34000 msec to the lefthand and from 35000 to the righthand. The predetermined time window preferably is 1000 msec.

As prosodic features, the slopes s_1 and s_2 of the obtained regression lines (in units of Hz/s) and/or the F_0 values v_1 , v_2 at the positions of f_1 and f_2 but estimated through the regression can be extracted. The advantage for using v_1 and v_2 instead of f_1 and f_2 is that determination of the F_0 values at f_1 (Offset) and f_2 (Onset) is rather faulty due to the transient behaviour of the spoken language at these boundaries. In accordance with f_1 and f_2 , also the difference $v_2 - v_1$ can be used as a prosodic feature.

Thereupon, the fundamental frequency itself can be used as prosodic feature. As mentioned above, it has been observed that the fundamental frequency in most cases declines continuously along a spoken sentence. It is also observed that nonvocal

sections in a speech stream correspond to gaps in the according fundamental frequency contour.

In a further embodiment at least two features are utilized, e.g. F0 (voiced-voiceless) and the audio signal intensity. Combining two features enhances robustness of the proposed mechanism. In addition to the predetermined time interval of 1000 msec, the robustness of the proposed mechanism can be enhanced by extracting corresponding features also within varied time intervals of e.g. 500, 2000, and 4000 msec and to compare the results.

The flow diagrams depicted in Figures 5a - 5c illustrate procedural steps for the segmentation of continuous speech in accordance with the invention.

In Fig. 5a, a digitized speech signal 600 is input to an F0 processor depicted in Fig. 3 that computes 610 a continuous F0 data from the speech signal. Only by the criterion F0 = ON/OFF, as described beforehand, the speech signal is presegmented 620 into speech segments. For each segment 630 it is evaluated 640 whether F0 is defined or no defined. In case of a not defined F0 (i.e. F0 = OFF) a candidate segment boundary is assumed as described above and, starting from that boundary, prosodic features be computed 650. The feature values are input into a classification tree (s. Fig. 6) and each candidate segment is classified thereby revealing, as a result, the existence or non-existence of a semantic or syntactic speech unit.

The 'Segmentation' step 620 in Fig. 5a is depicted in more detail in Fig. 5b. After initializing 700 variables "state", "start" and "stop", for each segment (frame) 710 it is checked whether F0 is defined (= ON) or not defined (= OFF) 720. In case of F0 = ON, it is further checked 730 whether the variable "state" is equal zero or not. If so, it is written 740 to the variables "state", "start" and "stop" whereby "state" is set '1', "start" is set 'stop + 1' and "stop" is set to the current value of "start" 750.

Thereafter it is continued with a new segment (frame).

In case step 720 reveals that F0 is not defined, i.e. F0 = OFF, it is also checked 770 whether variable "state" is '0'. If so, variable "stop" is set to '0' and thereafter continued 760 with a next frame. If variable "state" is not '0', it is written 790 to the three variables whereby variable "state" is set 1, "start" is set 'stop + 1' and "stop" is set to the current value of "start" 800.

The 'Compute features' step 650 shown in Fig. 5a is now depicted in more detail referring to Fig. 5c. In a first step 900, starting from a candidate boundary with F0 = OFF, F0 itself is used as prosodic features and computed accordingly. IN a further step 910, prosodic features in a time window lying before the candidate boundary are computed. In a next step 920, prosodic features are computed in a time window lying after the candidate boundary.

Fig. 6 shows an exemplary embodiment of a binary decision or classification tree which is trained by way of analyzing sample text bodies and which comprises the same prosodic features depicted in Figures 4a - 4c (with the only exception of feature mi). Starting from the root node 1000, at each node a feature is questioned and, depending on the obtained value, it is decided on which path it is continued. If, for example, at node 1010 the question $f1 < 126$? is answered with YES then it is continued on a left branch 1020 and in case of the answer NO on the right branch 1030. Reaching one of the end nodes 1040 - 1120, a decision is made on whether a semantic boundary has been found or not. In case a boundary has been found, an underlying edit decision list (EDL) can be updated accordingly or the audio stream be segmented at the boundaries thus revealing audio segments.

For the training of such a classification tree, a certain amount of training data in the same field as the intended application

need to be gathered beforehand. This comprises speech data and their textual representation. The latter allows for the determination of syntactic or semantic boundaries from punctuation etc. present in the text. With a system which allows for the linking of audio data with the corresponding reference text (e.g. precited European Patent Application X XXX XXX (docket no. DE9-1999-0053 of present applicant) the semantic or syntactic segment boundaries in the audio can be inferred from the corresponding boundaries in the text. Having trained the classification tree, the textual representation of audio data in the desired application is not needed.

It is noted again hereby that, for the approach according to the present invention, there is no need for a speech recognizer invoked in the segmentation process.

Although the invention is preferably applicable in the field of speech recognition or in the above described field of automatic generation of EDLs it is understood that it can advantageously be applied to other technical fields of audio processing or preprocessing.

C L A I M S

1. A method for the segmentation of an audio stream into semantic or syntactic units wherein the audio stream is provided in a digitized format, comprising the steps of:

determining a fundamental frequency for the digitized audio stream;

detecting changes of the fundamental frequency in the audio stream;

determining candidate boundaries for the semantic or syntactic units depending on the detected changes of the fundamental frequency;

extracting at least one prosodic feature in the neighborhood of the candidate boundaries;

determining boundaries for the semantic or syntactic units depending on the at least one prosodic feature.
2. Method according to claim 1, wherein providing a threshold value for the voicedness of the fundamental frequency estimates and determining whether the voicedness of fundamental frequency estimates is lower than the threshold value.
3. Method according to claim 2, wherein defining an index function for the fundamental frequency having a value = 0 if the voicedness of the fundamental frequency is lower than the threshold value and having a value = 1 if the voicedness of the fundamental frequency is higher than the threshold value.

4. Method according to claim 3, wherein extracting at least one prosodic feature in an environment of the audio stream where the value of the index function is equal 0.
5. Method according to claim 4, wherein the environment is a time period between 500 and 4000 milliseconds.
6. Method according to any of the preceding claims, wherein the at least one prosodic feature is represented by the fundamental frequency.
7. Method according to any of the preceding claims, wherein extracting at least two prosodic features and combining the at least two prosodic features.
8. Method according to any of the preceding claims, wherein at first detecting speech and non-speech segments in the digitized audio stream and performing the steps of claim 1 thereafter only for detected speech segments.
9. Method according to claim 8, wherein utilizing the signal energy or signal energy changes, respectively, in the audio stream.
10. Method according to any of the preceding claims, wherein performing a prosodic feature classification based on a predetermined classification tree.
11. Method according to any of claims 1 to 9, wherein a prosodic feature classification is performed by means of a neural net, an n-dimensional clustering or the like.
12. An article of manufacture comprising a computer usable medium having computer readable program code means

embodied therein for causing segmentation of an audio stream into semantic or syntactic units wherein the audio stream is provided in a digitized format, the computer readable program code means in the article of manufacture comprising computer readable program code means for causing a computer to effect:

determining a fundamental frequency for the digitized audio stream;

detecting changes of the fundamental frequency in the audio stream;

determining candidate boundaries for the semantic or syntactic units depending on the detected changes of the fundamental frequency;

extracting at least one prosodic feature in the neighborhood of the candidate boundaries;

determining boundaries for the semantic or syntactic units depending on the at least one prosodic feature.

13. Digital audio processing system for segmentation of a digitized audio stream into semantic or syntactic units comprising:

means for determining a fundamental frequency for the digitized audio stream,

means for detecting changes of the fundamental frequency in the audio stream,

means for determining candidate boundaries for the semantic or syntactic units depending on the detected changes of the fundamental frequency, and

means for extracting at least one prosodic feature in the neighborhood of the candidate boundaries.

means for determining boundaries for the semantic or syntactic units depending on the at least one prosodic feature.

14. Audio processing system according to claim 13, further comprising means for generating an index function for the voicedness of the fundamental frequency having a value = 0 if the voicedness of the fundamental frequency is lower than a predetermined threshold value and having a value = 1 if the voicedness fundamental frequency is higher than the threshold value.
15. Audio processing system according to claim 13 or 14, further comprising means for detecting speech and non-speech segments in the digitized audio stream, particularly for detecting and analyzing the signal energy or signal energy changes, respectively, in the audio stream.




FIG. 1




FIG. 4A




FIG. 2

2 / 7

3 / 7

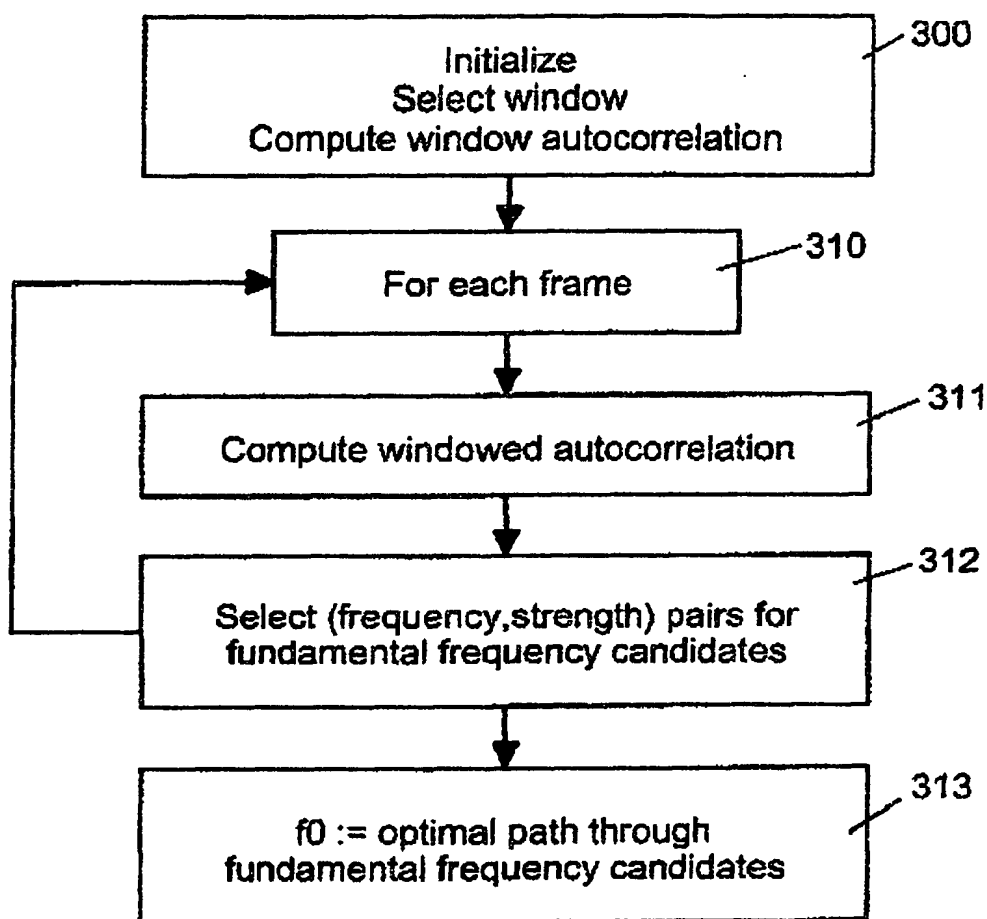


FIG. 3

4 / 7




FIG. 4B




FIG. 4C

5 / 7




FIG. 5A




FIG. 5C

6 / 7

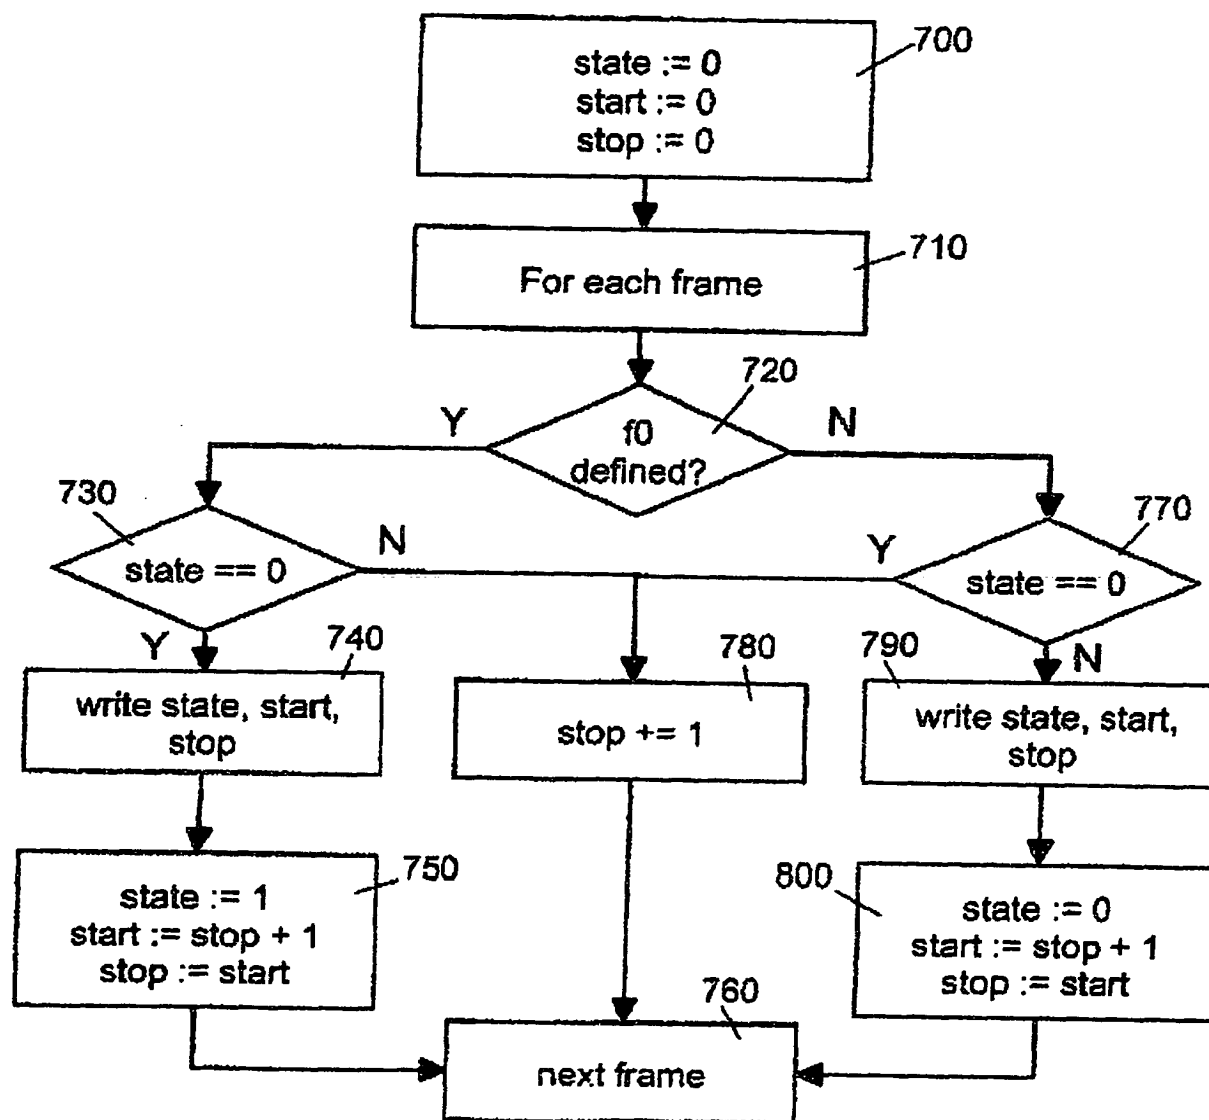


FIG.5B

7 / 7

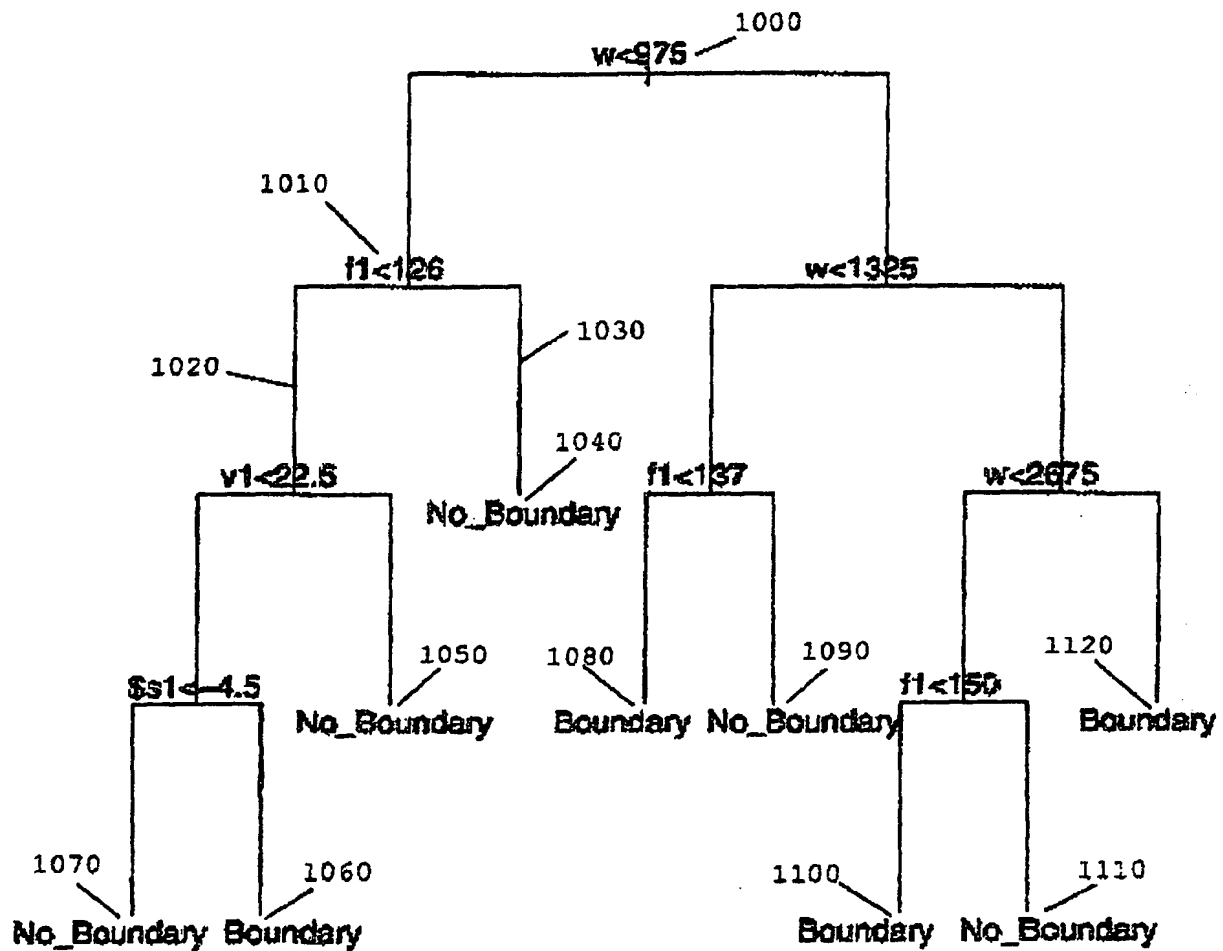


FIG. 6

A B S T R A C T

A digitized speech signal (600) is input to an F0 (fundamental frequency) processor that computes (610) a continuous F0 data from the speech signal. By the criterion voicing state transition (voiced/unvoiced transitions) the speech signal is presegmented (620) into segments. For each segment (630) it is evaluated (640) whether F0 is defined or not defined i.e. whether F0 is ON or OFF. In case of F0 = OFF a candidate segment boundary is assumed as described above and, starting from that boundary, prosodic features are computed (650). The feature values are input into a classification tree and each candidate segment is classified thereby revealing, as a result, the existence or non-existence of a semantic or syntactic speech unit.

(Fig. 6a)